Atlas Neural Networks: More Predictable Learning Through Fixed Context Routing

Key Observation

In a conventional ReLU neural network, presenting a training example fixes the activation pattern, so the forward and backward passes are locally linear. However, the weight update changes the activation boundaries themselves, meaning the routing for previously learned examples may also change.

Thus a single SGD update has two effects:

- **Linear:** modifies the parameters on the current computation path.
- **Nonlinear:** changes the computation paths taken by other examples.

This changing routing is a major source of interference between stored memories.

Atlas Perspective

In an Atlas neural network, routing is determined solely by an input-dependent context function

$$C(x)$$

rather than by trainable parameters.

The selected parameter set

$$P(C, x)$$

may depend on both the context and the current layer input, but the context itself is fixed for a given input.

Consequently:

- the computation graph for each training example remains fixed throughout training;
- parameter updates modify only the selected linear mappings;
- previously learned examples experience predictable linear perturbations rather than routing changes.

Deep networks can still amplify these perturbations through repeated matrix multiplication, but the source of nonlinearity arising from changing activation regions is removed.

Potential Advantages

Predictable interference

- Examples interfere only through shared parameter blocks.
- Context overlap provides a direct estimate of potential interference.

Targeted replay

- Replay only examples sharing modified context regions.
- Avoid unnecessary rehearsal of unaffected examples.

Larger learning rates

- Fixed routing removes one source of optimization instability.
- More aggressive optimization schedules may be possible.

Block-wise optimization

- Each context behaves approximately as a local linear model.
- Individual parameter blocks could use specialized optimizers or second-order methods.

Recursive least-squares training

Rather than treating every update as generic SGD:

- maintain an online linear estimator for each frequently used parameter block;
- exploit the locally linear nature of the selected mapping.

This suggests connections with adaptive filtering and recursive least squares.

Analytical influence estimates

Since routing is fixed:

- affected parameter blocks are known exactly;
- affected training examples are approximately known;
- memory interference becomes measurable rather than opaque.

Context refinement

Monitor context usage statistics.

If one context accumulates excessive conflicting gradients:

- subdivide that context;
- create additional parameter regions;
- reduce destructive interference.

This resembles splitting overloaded leaves in a decision tree.

Research Direction

Atlas networks may be better viewed not as conventional neural networks with a different activation function, but as **collections of overlapping adaptive linear systems whose composition is selected by input context**.

This viewpoint suggests borrowing optimization methods from adaptive signal processing, online least-squares estimation, Kalman filtering, and block-coordinate optimization rather than relying exclusively on standard stochastic gradient descent.

The principal hypothesis is that **fixed context routing makes learning dynamics substantially more predictable**, enabling more accurate estimates of interference, more targeted replay strategies, and potentially more efficient optimization algorithms.